Economics 203: Section 8

Jeffrey Naecker

March 4, 2014

1 Logistics

Don't be shy about coming to office hours! They are Thursday, 4-6 pm in Landau 360.

2 Concepts

2.1 Repeated Games

We consider the case where a so-called **stage game** is repeated some number of times. In this application, we have three important variables:

- The horizon: Finite or infinite
- The payoffs (if infinite): discounted, average payoff criterion, or overtaking criterion
- The solution concept: Nash equilibrium or subgame perfect Nash equilibrium

For infinitely repeated games, there are a few key points to remember:

- Playing a Nash equilibrium of the stage game in every period is a Nash equilibrium of the repeated game.
- When using discounted payoffs: Players can usually support playing cooperatively in SPNE by using a **Nash reversion** strategy, wherein if any player deviates from the cooperative outcome, play reverts to a NE of the stage game. Such a SPNE will usually only hold for certain δ. Note that for such constructions, you only need to check against deviations from the equilibrium path. Off-path, we are playing a stage-game NE each period, which by the above point is a NE of the subgame, and thus can be part of a SPNE for the repeated game as a whole.

- A common idiom you will use a lot when examining whether a deviation is profitable when using discounted payoffs:

$$\frac{1}{1-\delta}(\text{Equilibrium payoff}) \geq (\text{Deviation payoff}) + \frac{\delta}{1-\delta}(\text{Punishment payoff})$$

- Consider a strategy profile σ of the stage game. For each player i, we can define that player's **min-max payoff** $\pi_i^m = \min_{\sigma_{-i}} \max_{\sigma_i} \pi_i(\sigma)$.[1] Consider a payoff profile that is in the convex hull of stage game NE payoffs and also has no player receiving worse than his min-max payoff. Then the **folk theorem** says (when using average payoffs) that this payoff profile can be supported as the average payoffs for some NE of the repeated game. The NE folk theorem is similar if using discounted payoffs, except that δ has to be sufficiently close to unity.
- Note that punishment in NE can be minimax because the threat doesn't need to be credible. By contrast, in SPNE the worst we can credibly punish is the worst NE of stage game.

3 Problems

Problem 1. *Problem bank #82.*

Consider the following game in normal form (where player 1 chooses rows, player 2 chooses columns, and player 1's payoff is given first in each cell).

	a	b	c
A	5,3	5,5	3,4
B	4,10	9,9	4,11
C	3,3	11,4	5,5

(a) Suppose that this game is repeated a finite number of times, and that the discount factor (δ) is unity. Characterize the set of subgame perfect equilibria.

(b) Suppose that the game is repeated infinitely, and that δ is less than unity. Suppose that the players try to sustain the cooperative choice (B, b) in every period through Nash reversion. For what values of δ is this an equilibrium? Is it subgame perfect?

Problem 2. *Problem bank #85.*

[1] Intuitively, this is player $i's$ best response payoff when his opponents are "out to get him", in the sense that they pick strategies to force him to his worst possible best-response.

Consider the linear Cournot oligopoly model where price p is given as a linear function of total output Q, $p = a - bQ$, and where output is produced by each firm i at constant marginal cost c (same for all firms). Suppose there are N firms. Suppose also that the stage game will be repeated infinitely, and that firms discount future payoffs at the rate δ. Let δ_N be the minimum discount factor for which it is possible to sustain complete collusion in subgame perfect Nash equilibrium, using Nash reversion. Solve for δ_N as a function of N. How does it vary with N? Is this what you would expect? Why, or why not?

Problem 5. *Problem bank #91.*

Consider an infinitely repeated game with 2 firms. In each period, two firms simultaneously choose a characteristic of their products from a binary set $\{a, b\}$, if they choose the same characteristic, they split the market equally and each gets stage payoff of $\frac{\pi}{2}$, otherwise firm 1 always gets the entire market (so firm 1 gets stage payoff of π and firm 2 gets 0).

(a) What is the minmax payoff for each firm in the stage game?

(b) What is the highest payoff for each firm that can be an outcome of a Nash equilibrium of the infinitely repeated game?

4 Solutions

Solution 1.

(a) Note that the unique NE of the stage game is (C, c). Thus the only possible SPNE of the finitely repeated game is that 1 plays C in every stage and 2 place c in every stage.

(b) Nash reversion implies for following strategies. Players 1 and 2 start out by playing B and b, respectively. If the outcome is every previous period was (B, b), the players continue to play B and b. Otherwise, they play C and c in every period thereafter.

On the equilibrium path, the agents each get payoff 9 every period. Assuming 1 is playing B, 2 would like to deviate to c for a payoff of 11. Similarly, assuming 2 is playing b, 1 would like to deviate to playing C, also for a payoff of 11. If a player deviates, after that period they will receive the stage-NE payoff of 5 indefinitely. To sustain cooperation, we thus require that

$$\frac{1}{1-\delta}9 \geq 11 + \frac{\delta}{1-\delta}5,$$

which is equivalent to $\delta \geq \frac{1}{3}$.

This is indeed subgame perfect. Note that the condition above ensures that no player wants to deviate from the equilibrium path. Off the path, the specific strategies are NE of the stage game, and thus no player has incentive to deviate there either.

Solution 2.

Nash reversion specifies the following strategies in the repeated game: Each firm starts by producing the collusion quantity in the first repetition of the game. If in all previous period, all firms have played that quanitity, all firms continue to play that quantity in the current period. Otherwise, all firms play the Cournot quantity for all remaining repetitions of the stage game.

First, some notation: In complete collusion, all firms produce identical amounts such that the total quantity is the monopoly quantity, Q^M. Let the profits for each firm in this case be π^M. If one firm deviates, they will best-respond to all other firms producing q^M. That deviating firm will get profit π^D. From then on, all firms produce q^C, the Cournot quantity, for profits π^C.

For a single-shot deviation not to be profitable, we require the following: The profits from playing the collusion quantity forever must be greater than the profits from getting the deviation profit one period, followed by the Cournot profit in every following period. That is, we need

$$\frac{1}{1-\delta}\pi^M \geq \pi^D + \frac{\delta}{1-\delta}\pi^C.$$

It remains to find π^M, π^D, and π^C, and then solve for δ.

- First, we find the collusion profits. Monopoly quantity is given by $Q^M = \arg\max_Q (a - bQ - c)Q$, which we can solve to find $Q^M = \frac{a-c}{2b}$. Then each colluding firm produces $q^M = \frac{a-c}{2bN}$. Plugging in, we find price is given by $p^M = a - b\frac{a-c}{2b} = \frac{a+c}{2}$. Thus, collision profits are $\pi^M = (p^M - c)q^M = \left(\frac{a+c}{2} - c\right)\frac{a-c}{2bN}$. This gives $\pi^M = \frac{(a-c)^2}{4bN}$.
- In the Cournot equilibrium, all firms must be best-responding to the quantity produced by their competitors. Let q_i be firm i's quantity, and q_{-i} be the total quantity produced by that firm's competitors. Then $q_i^C = \arg\max_q \left[a - b(q + q_{-i}) - c\right] q$. The FOC give $a - 2bq - bq_{-i} - c = 0$. But we know that by all firms are producing the same quantity, so $q_{-i} = (N-1)q$. Thus we find $q^C = \frac{a-c}{b(N+1)}$. Then $p^C = a - bN\frac{a-c}{b(N+1)} = \frac{a+Nc}{N+1}$, and $\pi^C = (p^C - c)q^C = \frac{(a-c)^2}{b(N+1)^2}$.
- It remains to find the deviation profits. In this case, the deviating firm can assume that all other firms are producing the collusion amount, so that $q_{-i} = (N-1)q^m = \frac{N-1}{N}\frac{a-c}{2b} \equiv \bar{Q}$. The deviation quantity is given by $q^D = \arg\max_q \left[a - b(q + \bar{Q}) - c\right] q$. The FOC gives $q^D = \frac{a-c}{2b} - \frac{\bar{Q}}{2} = \frac{a-c}{2b} - \frac{N-1}{N}\frac{a-c}{4b} = \frac{a-c}{4b}\frac{N+1}{N}$. Thus $\pi^D = \left[a - b\left(\frac{a-c}{4b}\frac{N+1}{N} + \frac{N-1}{N}\frac{a-c}{2b}\right) - c\right]\frac{a-c}{4b}\frac{N+1}{N}$. With a bit of algebra we can reduce this to $\pi^D = \frac{(a-c)^2(N+1)^2}{16bN^2}$.

Thus our Nash revision condition becomes

$$\frac{1}{1-\delta}\frac{(a-c)^2}{4bN} \geq \frac{(a-c)^2(N+1)^2}{16bN^2} + \frac{\delta}{1-\delta}\frac{(a-c)^2}{b(N+1)^2}.$$

To solve this, I recommend multiply by $\frac{4bN(1-\delta)}{(a-c)^2}$. Then we get

$$1 \geq (1-\delta)\frac{(N+1)^2}{4N} + \delta\frac{4N}{(N+1)^2}.$$

Let's do a change of variables, setting $x = \frac{(N+1)^2}{4N}$. Then our condition reduces further to $1 \geq (1-\delta)x + \delta\frac{1}{x}$. We can solve this to find that $\delta \geq \frac{x}{1+x}$. Substituting back in, we finally find

$$\delta \geq \frac{\frac{(N+1)^2}{4N}}{1+\frac{(N+1)^2}{4N}} = \frac{(N+1)^2}{4N+(N+1)^2} \equiv \delta_N.$$

Note that $\delta_N \to 1$ as $N \to \infty$. This means that collusion requires firms to be more and more patient as the number of firms grows. To see why, note that as $N \to \infty$, we have $\pi^M \to 0$ and $\pi^C \to 0$ but $\pi^D \to \frac{(a-c)^2}{16b}$. This means that in the limit, firms get positive profit for deviating and infinitesimal profit for not deviating. So, firms must essentially not discount at all for collusion to be sustained.

Solution 3.

(a) Note that the normal form of the stage game is given by the following matrix:

	a	b
a	$\left(\frac{\pi}{2},\frac{\pi}{2}\right)$	$(\pi,0)$
b	$(\pi,0)$	$\left(\frac{\pi}{2},\frac{\pi}{2}\right)$

Suppose that player firm 1's strategy is to mix with probability δ_1, while firm 2 mixes with probability δ_2.

- Firm 1's min-max payoff is given by

$$\min_{\delta_2}\max_{\delta_1}\left(\delta_1\delta_2\frac{\pi}{2} + \delta_1(1-\delta_2)\pi + (1-\delta_1)\delta_2\pi + (1-\delta_1)(1-\delta_2)\frac{\pi}{2}\right).$$

Note that the main expression can be re-written as $\frac{\pi}{2} + \delta_1(\frac{1}{2}-\delta_2)\pi + \delta_2\frac{\pi}{2}$. First, let δ_2 be given. Then only the middle term of this expression matters for our maximization. In particular, we set $\delta_1 = 1$ if $\delta_2 < \frac{1}{2}$, for a payoff of $\pi - \delta_2\frac{\pi}{2}$, and $\delta_1 = 0$ if $\delta_2 > \frac{1}{2}$, for a payoff of $\frac{\pi}{2} + \delta_2\frac{\pi}{2}$. If $\delta_2 = \frac{1}{2}$ then any $\delta_1 \in [0,1]$ give a payoff of $\frac{3}{4}\pi$.

Taking this maximization under δ_1 into account, we can now reduce the overall problem to

$$\min_{\delta_2}\begin{cases} \pi - \delta_2\frac{\pi}{2} & \text{if } \delta_2 < \frac{1}{2} \\ \frac{\pi}{2} + \delta_2\frac{\pi}{2} & \text{if } \delta_2 > \frac{1}{2} \\ \frac{3}{4}\pi & \text{if } \delta_2 = \frac{1}{2}. \end{cases}$$

This problem is solved by $\delta_2 = \frac{1}{2}$. Thus firm 1's min-max payoffs are $\frac{3}{4}\pi$.

- Firm 2's min-max payoffs are given by

$$\min_{\delta_1} \max_{\delta_2} \left(\delta_1 \delta_2 \frac{\pi}{2} + (1 - \delta_1)(1 - \delta_2) \frac{\pi}{2} \right).$$

Again, note this expression can be re-written as $\frac{\pi}{2} + \delta_2 (2\delta_1 - 1) \frac{\pi}{2} - \delta_1 \frac{\pi}{2}$. If $\delta_1 > \frac{1}{2}$ we set $\delta_2 = 1$ for payoff of $\delta_1 \frac{\pi}{2}$. If $\delta_1 < \frac{1}{2}$ we set $\delta_2 = 0$ for payoff of $\frac{\pi}{2} - \delta_1 \frac{\pi}{2}$. And if $\delta_1 = \frac{1}{2}$ then any $\delta_2 \in [0, 1]$ gives a payoff of $\frac{1}{4}\pi$.

Taking this maximization under δ_2 into account, we now solve

$$\min_{\delta_1} \begin{cases} \delta_1 \frac{\pi}{2} & \text{if } \delta_1 > \frac{1}{2} \\ \frac{\pi}{2} - \delta_1 \frac{\pi}{2} & \text{if } \delta_1 < \frac{1}{2} \\ \frac{1}{4}\pi & \text{if } \delta_1 = \frac{1}{2}. \end{cases}$$

This problem is solved by $\delta_1 = \frac{1}{2}$. Thus firm 2's min-max payoffs are $\frac{1}{4}\pi$.

(b) First, recall that the folk theorem says that in any NE of the repeated game, each player much receive average payoffs at least as great as their min-max payoffs. But since the game is constant-sum, if one firm were to get strictly more than their min-max payoffs, the other firm would necessarily get strictly less. Thus the firms must each get exactly their min-max payoffs in any NE of the infinitely repeated game.